

Amendments to the Specification:

Please replace the paragraph on page 5, lines 3-22 with the following amended paragraph:

The present invention provides a virtual stored data management system. In one embodiment, the management system includes one or more hosts and a plurality of data storage elements functionally coupled to the hosts. Each data storage element includes a host network attachment, data transfer means, a storage controller, and permanent data storage media. The permanent data storage media is organized with management information uniquely associated with units of the data such that the management information may be manipulated at nodes that are in several different locations within the management system substantially simultaneously. Thus, the organization of the management processes allows for the management information to be processed, used, changed, or modified at nodes that are in several different locations within the management system at any particular instance. Provision is made for the internal processes to discover the current location of the processing, for the location to be changed as directed, and for the processing to be kept consistent when done in more than one place simultaneously.

Please replace the paragraph extending from page 25, line 1 to page 26, line 18 with the following amended paragraph:

With reference to **Figure 9**, a block diagram illustrating a storage virtualization system is depicted in accordance with a preferred embodiment of the present invention. System **900** includes similar components to system **800** including hosts **904-908**, network **902**, storage server **910**, storage controller **912**, and storage **914**. However, to solve the problem described above of the overload of key processing units or transfer paths, the level 1, 2, and 3 processing is moved at the request of a host or as a consequence of internal processes that note the contention to more strategic locations. Since Host 1 **904** and Host i **906** have access to the network

902, moving level 1 and level 2 resolution for devices 1 and replicating level 1 and level 2 resolutions for device 2 to the network 902 from the server 910 (which is actually using the processor in Host k 908) will significantly relieve the load on the server 910. Also moving the level 1 resolution for device 3 from the server 910 to the controller 914 will allow the transfer for device 3 to Host k 908 to go directly rather than through the server 910. The level 1 resolution for device 2 is also maintained in the server 910 for requests that do not go through the network 902. When Host 1 904 does data transfer with device 1 or device 2, the processing of level 1 and level 2 is done at a node in the network 902, the processing of the level 3 is passed through the server 910 to a node in the controller 912 and data flows from a node in storage 914 through the server 910 and the network 902 to the Host 1 904. When Host i 906 does transfer with device 1 or 2 it follows the same process. When Host i 906 accesses device 2 but finds the network busy or when host k 908 wishes to initiate transfer to device 2, the level 1 and level 2 resolution is processed at a node in the server 910 and communication is made with the network 902 to keep the processing of level 1 and level 2 for device 2 consistent between the two locations. When host i 906 wishes to initiate data transfer with device 3, it does so using whichever path to server 910 is less busy at the moment. The processing of level 1, level 2, and level 3 are all executed at a node in the controller 912 and transfer is initiated through the server 910 with Host i 906. When Host k 908 initiates transfer with device 3, the request is sent directly to a node in the controller 912 and transfer is initiated from device 3 through storage 914 and then directly between the controller 912 and Host k 908.